
A Residual Bootstrap for High-Dimensional Regression with Near Low-Rank Designs

Miles E. Lopes

Department of Statistics
University of California, Berkeley
Berkeley, CA 94720
mlopes@stat.berkeley.edu

Abstract

We study the residual bootstrap (RB) method in the context of high-dimensional linear regression. Specifically, we analyze the distributional approximation of linear contrasts $c^\top(\hat{\beta}_\rho - \beta)$, where $\hat{\beta}_\rho$ is a ridge-regression estimator. When regression coefficients are estimated via least squares, classical results show that RB consistently approximates the laws of contrasts, provided that $p \ll n$, where the design matrix is of size $n \times p$. Up to now, relatively little work has considered how additional structure in the linear model may extend the validity of RB to the setting where $p/n \asymp 1$. In this setting, we propose a version of RB that resamples residuals obtained from ridge regression. Our main structural assumption on the design matrix is that it is nearly low rank — in the sense that its singular values decay according to a power-law profile. Under a few extra technical assumptions, we derive a simple criterion for ensuring that RB consistently approximates the law of a given contrast. We then specialize this result to study confidence intervals for mean response values $X_i^\top \beta$, where X_i^\top is the i th row of the design. More precisely, we show that conditionally on a Gaussian design with near low-rank structure, RB *simultaneously* approximates all of the laws $X_i^\top(\hat{\beta}_\rho - \beta)$, $i = 1, \dots, n$. This result is also notable as it imposes no sparsity assumptions on β . Furthermore, since our consistency results are formulated in terms of the Mallows (Kantorovich) metric, the existence of a limiting distribution is not required.

1 Introduction

Until recently, much of the emphasis in the theory of high-dimensional statistics has been on “first order” problems, such as estimation and prediction. As the understanding of these problems has become more complete, attention has begun to shift increasingly towards “second order” problems, dealing with hypothesis tests, confidence intervals, and uncertainty quantification [1, 2, 3, 4, 5, 6]. In this direction, much less is understood about the effects of structure, regularization, and dimensionality — leaving many questions open. One collection of such questions that has attracted growing interest deals with the operating characteristics of the *bootstrap* in high dimensions [7, 8, 9]. Due to the fact that bootstrap is among the most widely used tools for approximating the sampling distributions of test statistics and estimators, there is much practical value in understanding what factors allow for the bootstrap to succeed in the high-dimensional regime.

The regression model and linear contrasts. In this paper, we focus our attention on high-dimensional linear regression, and our aim is to know when the residual bootstrap (RB) method consistently approximates the laws of *linear contrasts*. (A review of RB is given in Section 2.)

To specify the model, suppose that we observe a response vector $Y \in \mathbb{R}^n$, generated according to

$$Y = X\beta + \varepsilon, \quad (1)$$

where $X \in \mathbb{R}^{n \times p}$ is the observed design matrix, $\beta \in \mathbb{R}^p$ is an unknown vector of coefficients, and the error variables $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)$ are drawn i.i.d. from an unknown distribution F_0 , with mean 0 and unknown variance $\sigma^2 < \infty$. As is conventional in high-dimensional statistics, we assume the model (1) is embedded in a sequence of models indexed by n . Hence, we allow X , β , and p to vary implicitly with n . We will leave p/n unconstrained until Section 3.3, where we will assume $p/n \asymp 1$ in Theorem 3, and then in Section 3.4, we will assume further that p/n is bounded strictly between 0 and 1. The distribution F_0 is fixed with respect to n , and none of our results require F_0 to have more than four moments.

Although we are primarily interested in cases where the design matrix X is deterministic, we will also study the performance of the bootstrap conditionally on a Gaussian design. For this reason, we will use the symbol $\mathbb{E}[\dots|X]$ even when the design is non-random so that confusion does not arise in relating different sections of the paper. Likewise, the symbol $\mathbb{E}[\dots]$ refers to unconditional expectation over all sources of randomness. Whenever the design is random, we will assume $X \perp\!\!\!\perp \varepsilon$, denoting the distribution of X by \mathbb{P}_X , and the distribution of ε by \mathbb{P}_ε .

Within the context of the regression, we will be focused on linear contrasts $c^\top(\hat{\beta} - \beta)$, where $c \in \mathbb{R}^p$ is a fixed vector and $\hat{\beta} \in \mathbb{R}^p$ is an estimate of β . The importance of contrasts arises from the fact that they unify many questions about a linear model. For instance, testing the significance of the i th coefficient β_i may be addressed by choosing c to be the standard basis vector $c^\top = e_i^\top$. Another important problem is quantifying the uncertainty of point predictions, which may be addressed by choosing $c^\top = X_i^\top$, i.e. the i th row of the design matrix. In this case, an approximation to the law of the contrast leads to a confidence interval for the mean response value $\mathbb{E}[Y_i] = X_i^\top \beta$. Further applications of contrasts occur in the broad topic of ANOVA [10].

Intuition for structure and regularization in RB. The following two paragraphs explain the core conceptual aspects of the paper. To understand the role of regularization in applying RB to high-dimensional regression, it is helpful to think of RB in terms of two ideas. First, if $\hat{\beta}_{\text{LS}}$ denotes the ordinary least squares estimator, then it is a simple but important fact that contrasts can be written as $c^\top(\hat{\beta}_{\text{LS}} - \beta) = a^\top \varepsilon$ where $a^\top := c^\top(X^\top X)^{-1}X^\top$. Hence, if it were possible to sample directly from F_0 , then the law of any such contrast could be easily determined. Since F_0 is unknown, the second key idea is to use the residuals of *some* estimator $\hat{\beta}$ as a proxy for samples from F_0 . When $p \ll n$, the least-squares residuals are a good proxy [11, 12]. However, it is well-known that least-squares tends to overfit when $p/n \asymp 1$. When $\hat{\beta}_{\text{LS}}$ fits “too well”, this means that its residuals are “too small”, and hence they give a poor proxy for F_0 . Therefore, by using a regularized estimator $\hat{\beta}$, overfitting can be avoided, and the residuals of $\hat{\beta}$ may offer a better way of obtaining “approximate samples” from F_0 .

The form of regularized regression we will focus on is *ridge regression*:

$$\hat{\beta}_\rho := (X^\top X + \rho I_{p \times p})^{-1} X^\top Y, \quad (2)$$

where $\rho > 0$ is a user-specified regularization parameter. As will be seen in Sections 3.2 and 3.3, the residuals obtained from ridge regression lead to a particularly good approximation of F_0 when the design matrix X is nearly low-rank, in the sense that most of its singular values are close to 0. In essence, this condition is a form of sparsity, since it implies that the rows of X nearly lie in a low-dimensional subspace of \mathbb{R}^p . However, this type of structural condition has a significant advantage over the more well-studied assumption that β is sparse. Namely, the assumption that X is nearly low-rank can be inspected directly in practice — whereas sparsity in β is typically unverifiable. In fact, our results will impose no conditions on β , other than that $\|\beta\|_2$ remains bounded as $(n, p) \rightarrow \infty$. Finally, it is worth noting that the occurrence of near low-rank design matrices is actually very common in applications, and is often referred to as *collinearity* [13, ch. 17].

Contributions and outline. The primary contribution of this paper is a complement to the work of Bickel and Freedman [12] (hereafter B&F 1983) — who showed that in general, the RB method fails

to approximate the laws of least-squares contrasts $c^\top(\hat{\beta}_{\text{LS}} - \beta)$ when $p/n \asymp 1$. Instead, we develop an alternative set of results, proving that even when $p/n \asymp 1$, RB can successfully approximate the laws of “ridged contrasts” $c^\top(\hat{\beta}_\rho - \beta)$ for many choices of $c \in \mathbb{R}^p$, provided that the design matrix X is nearly low rank. A particularly interesting consequence of our work is that RB successfully approximates the law $c^\top(\hat{\beta}_\rho - \beta)$ for a certain choice of c that was shown in B&F 1983 to “break” RB when applied to least-squares. Specifically, such a c can be chosen as one of the rows of X with a high *leverage score* (see Section 4). This example corresponds to the practical problem of setting confidence intervals for mean response values $\mathbb{E}[Y_i] = X_i^\top \beta$. (See [12, p. 41], as well as Lemma 2 and Theorem 4 in Section 3.4). Lastly, from a technical point of view, a third notable aspect of our results is that they are formulated in terms of the Mallows- ℓ_2 metric, which frees us from having to impose conditions that force a limiting distribution to exist.

Apart from B&F 1983, the most closely related works we are aware of are the recent papers [7] and [8], which also consider RB in the high-dimensional setting. However, these works focus on role of sparsity in β and do not make use of low-rank structure in the design, whereas our work deals only with structure in the design and imposes no sparsity assumptions on β .

The remainder of the paper is organized as follows. In Section 2, we formulate the problem of approximating the laws of contrasts, and describe our proposed methodology for RB based on ridge regression. Then, in Section 3 we state several results that lay the groundwork for Theorem 4, which shows that that RB can successfully approximate all of the laws $\mathcal{L}(X_i^\top(\hat{\beta}_\rho - \beta)|X)$, $i = 1, \dots, n$, conditionally on a Gaussian design. Due to space constraints, all proofs are deferred to material that will appear in a separate work.

Notation and conventions. If U and V are random variables, then $\mathcal{L}(U|V)$ denotes the law of U , conditionally on V . If a_n and b_n are two sequences of real numbers, then the notation $a_n \lesssim b_n$ means that there is an absolute constant $\kappa_0 > 0$ and an integer $n_0 \geq 1$ such that $a_n \leq \kappa_0 b_n$ for all $n \geq n_0$. The notation $a_n \asymp b_n$ means that $a_n \lesssim b_n$ and $b_n \lesssim a_n$. For a square matrix $A \in \mathbb{R}^{k \times k}$ whose eigenvalues are real, we denote them by $\lambda_{\min}(A) = \lambda_k(A) \leq \dots \leq \lambda_1(A) = \lambda_{\max}(A)$.

2 Problem setup and methodology

Problem setup. For any $c \in \mathbb{R}^p$, it is clear that conditionally on X , the law of $c^\top(\hat{\beta}_\rho - \beta)$ is completely determined by F_0 , and hence it makes sense to use the notation

$$\Psi_\rho(F_0; c) := \mathcal{L}(c^\top(\hat{\beta}_\rho - \beta) \mid X). \quad (3)$$

The problem we aim to solve is to approximate the distribution $\Psi_\rho(F_0; c)$ for suitable choices of c .

Review of the residual bootstrap (RB) procedure. We briefly explain the steps involved in the residual bootstrap procedure, applied to the ridge estimator $\hat{\beta}_\rho$ of β . To proceed somewhat indirectly, consider the following “bias-variance” decomposition of $\Psi_\rho(F_0; c)$, conditionally on X ,

$$\Psi_\rho(F_0; c) = \underbrace{\mathcal{L}(c^\top(\hat{\beta}_\rho - \mathbb{E}[\hat{\beta}_\rho|X]) \mid X)}_{=: \Phi_\rho(F_0; c)} + \underbrace{c^\top(\mathbb{E}[\hat{\beta}_\rho|X] - \beta)}_{=: \text{bias}(\Phi_\rho(F_0; c))}. \quad (4)$$

Note that the distribution $\Phi(F_0; c)$ has mean zero, and so that the second term on the right side is the bias of $\Phi_\rho(F_0; c)$ as an estimator of $\Psi_\rho(F_0; c)$. Furthermore, the distribution $\Phi_\rho(F_0; c)$ may be viewed as the “variance component” of $\Psi_\rho(F_0; c)$. We will be interested in situations where the regularization parameter ρ may be chosen small enough so that the bias component is small. In this case, one has $\Psi_\rho(F_0; c) \approx \Phi_\rho(F_0; c)$, and then it is enough to find an approximation to the law $\Phi_\rho(F_0; c)$, which is unknown. To this end, a simple manipulation of $c^\top(\hat{\beta}_\rho - \mathbb{E}[\hat{\beta}_\rho])$ leads to

$$\Phi_\rho(F_0; c) = \mathcal{L}(c^\top(X^\top X + \rho I_{p \times p})^{-1} X^\top \varepsilon \mid X). \quad (5)$$

Now, to approximate $\Phi_\rho(F_0; c)$, let \hat{F} be any centered estimate of F_0 . (Typically, \hat{F} is obtained by using the centered residuals of some estimator of β , but this is not necessary in general.) Also, let $\varepsilon^* = (\varepsilon_1^*, \dots, \varepsilon_n^*) \in \mathbb{R}^n$ be an i.i.d. sample from \hat{F} . Then, replacing ε with ε^* in line (5) yields

$$\Phi_\rho(\hat{F}; c) = \mathcal{L}(c^\top(X^\top X + \rho I_{p \times p})^{-1} X^\top \varepsilon^* \mid X). \quad (6)$$

At this point, we define the (random) measure $\Phi_\rho(\hat{F}; c)$ to be the RB approximation to $\Phi_\rho(F_0; c)$. Hence, it is clear that the RB approximation is simply a “plug-in rule”.

A two-stage approach. An important feature of the procedure just described is that we are free to use any centered estimator \hat{F} of F_0 . This fact offers substantial flexibility in approximating $\Psi_\rho(F_0; c)$. One way of exploiting this flexibility is to consider a two-stage approach, where a “pilot” ridge estimator $\hat{\beta}_\varrho$ is used to first compute residuals whose centered empirical distribution function is \hat{F}_ϱ , say. Then, in the second stage, the distribution \hat{F}_ϱ is used to approximate $\Phi_\rho(F_0; c)$ via the relation (6). To be more detailed, if $(\hat{e}_1(\varrho), \dots, \hat{e}_n(\varrho)) = \hat{e}(\varrho) := Y - X\hat{\beta}_\varrho$ are the residuals of $\hat{\beta}_\varrho$, then we define \hat{F}_ϱ to be the distribution that places mass $1/n$ at each of the values $\hat{e}_i(\varrho) - \bar{e}(\varrho)$ with $\bar{e}(\varrho) := \frac{1}{n} \sum_{i=1}^n \hat{e}_i(\varrho)$. Here, it is important to note that the value ϱ is chosen to optimize \hat{F}_ϱ as an approximation to F_0 . By contrast, the choice of ρ depends on the relative importance of width and coverage probability for confidence intervals based on $\Phi_\rho(\hat{F}_\varrho; c)$. Theorems 1, 3, and 4 will offer some guidance in selecting ϱ and ρ .

Resampling algorithm. To summarize the discussion above, if B is user-specified number of bootstrap replicates, our proposed method for approximating $\Psi_\rho(F_0; c)$ is given below.

1. Select ρ and ϱ , and compute the residuals $\hat{e}(\varrho) = Y - X\hat{\beta}_\varrho$.
2. Compute the centered distribution function \hat{F}_ϱ , putting mass $1/n$ at each $\hat{e}_i(\varrho) - \bar{e}(\varrho)$.
3. For $j = 1, \dots, B$:
 - Draw a vector $\varepsilon^* \in \mathbb{R}^n$ of n i.i.d. samples from \hat{F}_ϱ .
 - Compute $z_j := c^\top (X^\top X + \rho I_{p \times p})^{-1} X^\top \varepsilon^*$.
4. Return the empirical distribution of z_1, \dots, z_B .

Clearly, as $B \rightarrow \infty$, the empirical distribution of z_1, \dots, z_B converges weakly to $\Phi_\rho(\hat{F}_\varrho; c)$, with probability 1. As is conventional, our theoretical analysis in the next section will ignore Monte Carlo issues, and address only the performance of $\Phi_\rho(\hat{F}_\varrho; c)$ as an approximation to $\Psi_\rho(F_0; c)$.

3 Main results

The following metric will be central to our theoretical results, and has been a standard tool in the analysis of the bootstrap, beginning with the work of Bickel and Freedman [14].

The Mallows (Kantorovich) metric. For two random vectors U and V in a Euclidean space, the Mallows- ℓ_2 metric is defined by

$$d_2^2(\mathcal{L}(U), \mathcal{L}(V)) := \inf_{\pi \in \Pi} \left\{ \mathbb{E} \left[\|U - V\|_2^2 \right] : (U, V) \sim \pi \right\} \quad (7)$$

where the infimum is over the class Π of joint distributions π whose marginals are $\mathcal{L}(U)$ and $\mathcal{L}(V)$. It is worth noting that convergence in d_2 is strictly stronger than weak convergence, since it also requires convergence of second moments. Additional details may be found in the paper [14].

3.1 A bias-variance decomposition for bootstrap approximation

To give some notation for analyzing the bias-variance decomposition of $\Psi_\rho(F_0; c)$ in line (4), we define the following quantities based upon the ridge estimator $\hat{\beta}_\rho$. Namely, the variance is

$$v_\rho = v_\rho(X; c) := \text{var}(\Psi_\rho(F_0; c) | X) = \sigma^2 \|c^\top (X^\top X + \rho I_{p \times p})^{-1} X^\top\|_2^2.$$

To express the bias of $\Phi_\rho(F_0; c)$, we define the vector $\delta(X) \in \mathbb{R}^p$ according to

$$\delta(X) := \beta - \mathbb{E}[\hat{\beta}_\rho] = [I_{p \times p} - (X^\top X + \rho I_{p \times p})^{-1} X^\top X] \beta, \quad (8)$$

and then put

$$b_\rho^2 = b_\rho^2(X; c) := \text{bias}^2(\Phi_\rho(F_0; c)) = (c^\top \delta(X))^2. \quad (9)$$

We will sometimes omit the arguments of v_ρ and b_ρ^2 to lighten notation. Note that $v_\rho(X; c)$ does not depend on β , and $b_\rho^2(X; c)$ only depends on β through $\delta(X)$.

The following result gives a regularized and high-dimensional extension of some lemmas in Freedman's early work [11] on RB for least squares. The result does not require any structural conditions on the design matrix, or on the true parameter β .

Theorem 1 (consistency criterion). *Suppose $X \in \mathbb{R}^{n \times p}$ is fixed. Let \hat{F} be any estimator of F_0 , and let $c \in \mathbb{R}^p$ be any vector such that $v_\rho = v_\rho(X; c) \neq 0$. Then with \mathbb{P}_ε -probability 1, the following inequality holds for every $n \geq 1$, and every $\rho > 0$,*

$$d_2^2\left(\frac{1}{\sqrt{v_\rho}}\Psi_\rho(F_0; c), \frac{1}{\sqrt{v_\rho}}\Phi_\rho(\hat{F}; c)\right) \leq \frac{1}{\sigma^2}d_2^2(F_0, \hat{F}) + \frac{b_\rho^2}{v_\rho}. \quad (10)$$

Remarks. Observe that the normalization $1/\sqrt{v_\rho}$ ensures that the bound is non-trivial, since the distribution $\Psi_\rho(F_0; c)/\sqrt{v_\rho}$ has variance equal to 1 for all n (and hence does not become degenerate for large n). To consider the choice of ρ , it is simple to verify that the ratio b_ρ^2/v_ρ decreases monotonically as ρ decreases. Note also that as ρ becomes small, the variance v_ρ becomes large, and likewise, confidence intervals based on $\Phi_\rho(\hat{F}; c)$ become wider. In other words, there is a trade-off between the width of the confidence interval and the size of the bound (10).

Sufficient conditions for consistency of RB. An important practical aspect of Theorem 1 is that for any given contrast c , the variance $v_\rho(X; c)$ can be easily estimated, since it only requires an estimate of σ^2 , which can be obtained from \hat{F} . Consequently, whenever theoretical bounds on $d_2^2(F_0, \hat{F})$ and $b_\rho^2(X; c)$ are available, the right side of line (10) can be controlled. In this way, Theorem 1 offers a simple route for guaranteeing that RB is consistent. In Sections 3.2 and 3.3 to follow, we derive a bound on $\mathbb{E}[d_2^2(F_0, \hat{F})|X]$ in the case where \hat{F} is chosen to be \hat{F}_ℓ . Later on in Section 3.4, we study RB consistency in the context of prediction with a Gaussian design, and there we derive high probability bounds on both $v_\rho(X; c)$ and $b_\rho^2(X; c)$ where c is a particular row of X .

3.2 A link between bootstrap consistency and MSPE

If $\hat{\beta}$ is an estimator of β , its mean-squared prediction error (MSPE), conditionally on X , is defined as

$$\text{mspe}(\hat{\beta}|X) := \frac{1}{n}\mathbb{E}[\|X(\hat{\beta} - \beta)\|_2^2 | X]. \quad (11)$$

The previous subsection showed that in-law approximation of contrasts is closely tied to the approximation of F_0 . We now take a second step of showing that if the centered residuals of an estimator $\hat{\beta}$ are used to approximate F_0 , then the quality of this approximation can be bounded naturally in terms of $\text{mspe}(\hat{\beta}|X)$. This result applies to any estimator $\hat{\beta}$ computed from the observations (1).

Theorem 2. *Suppose $X \in \mathbb{R}^{n \times p}$ is fixed. Let $\hat{\beta}$ be any estimator of β , and let \hat{F} be the empirical distribution of the centered residuals of $\hat{\beta}$. Also, let F_n denote the empirical distribution of n i.i.d. samples from F_0 . Then for every $n \geq 1$,*

$$\mathbb{E}[d_2^2(\hat{F}, F_0) | X] \leq 2 \text{mspe}(\hat{\beta}|X) + 2 \mathbb{E}[d_2^2(F_n, F_0)] + \frac{2\sigma^2}{n}. \quad (12)$$

Remarks. As we will see in the next section, the MSPE of ridge regression can be bounded in a sharp way when the design matrix is approximately low rank, and there we will analyze $\text{mspe}(\hat{\beta}_\ell|X)$ for the pilot estimator. Consequently, when near low-rank structure is available, the only remaining issue in controlling the right side of line (12) is to bound the quantity $\mathbb{E}[d_2^2(F_n, F_0)|X]$. The very recent work of Bobkov and Ledoux [15] provides an in-depth study of this question, and they derive a variety bounds under different tail conditions on F_0 . We summarize one of their results below.

Lemma 1 (Bobkov and Ledoux, 2014). *If F_0 has a finite fourth moment, then*

$$\mathbb{E}[d_2^2(F_n, F_0)] \lesssim \log(n)n^{-1/2}. \quad (13)$$

Remarks. The fact that the *squared* distance is bounded at the rate of $\log(n)n^{-1/2}$ is an indication that d_2 is a rather strong metric on distributions. For a detailed discussion of this result, see Corollaries 7.17 and 7.18 in the paper [15]. Although it is possible to obtain faster rates when more stringent tail conditions are placed on F_0 , we will only need a fourth moment, since the $\text{mspe}(\widehat{\beta}|X)$ term in Theorem 2 will often have a slower rate than $\log(n)n^{-1/2}$, as discussed in the next section.

3.3 Consistency of ridge regression in MSPE for near low rank designs

In this subsection, we show that when the tuning parameter ϱ is set at a suitable rate, the pilot ridge estimator $\widehat{\beta}_\varrho$ is consistent in MSPE when the design matrix is near low-rank — even when p/n is large, and without any sparsity constraints on β . We now state some assumptions.

A1. There is a number $\nu > 0$, and absolute constants $\kappa_1, \kappa_2 > 0$, such that

$$\kappa_1 i^{-\nu} \leq \lambda_i(\widehat{\Sigma}) \leq \kappa_2 i^{-\nu} \quad \text{for all } i = 1, \dots, n \wedge p.$$

A2. There are absolute constants $\theta, \gamma > 0$, such that for every $n \geq 1$, $\frac{\varrho}{n} = n^{-\theta}$ and $\frac{\rho}{n} = n^{-\gamma}$.

A3. The vector $\beta \in \mathbb{R}^p$ satisfies $\|\beta\|_2 \lesssim 1$.

Due to Theorem 2, the following bound shows that the residuals of $\widehat{\beta}_\varrho$ may be used to extract a consistent approximation to F_0 . Two other notable features of the bound are that it is *non-asymptotic* and *dimension-free*.

Theorem 3. Suppose that $X \in \mathbb{R}^{n \times p}$ is fixed and that Assumptions 1–3 hold, with $p/n \asymp 1$. Assume further that θ is chosen as $\theta = \frac{2\nu}{3}$ when $\nu \in (0, \frac{1}{2})$, and $\theta = \frac{\nu}{\nu+1}$ when $\nu > \frac{1}{2}$. Then,

$$\text{mspe}(\widehat{\beta}_\varrho|X) \lesssim \begin{cases} n^{-\frac{2\nu}{3}} & \text{if } \nu \in (0, \frac{1}{2}), \\ n^{-\frac{\nu}{\nu+1}} & \text{if } \nu > \frac{1}{2}. \end{cases} \quad (14)$$

Also, both bounds in (14) are tight in the sense that β can be chosen so that $\widehat{\beta}_\varrho$ attains either rate.

Remarks. Since the eigenvalues $\lambda_i(\widehat{\Sigma})$ are observable, they may be used to estimate ν and guide the selection of $\varrho/n = n^{-\theta}$. However, from a practical point of view, we found it easier to select ϱ via cross-validation in numerical experiments, rather than via an estimate of ν .

A link with Pinsker’s Theorem. In the particular case when F_0 is a centered Gaussian distribution, the “prediction problem” of estimating $X\beta$ is very similar to estimating the mean parameters of a Gaussian sequence model, with error measured in the ℓ_2 norm. In the alternative sequence-model format, the decay condition on the eigenvalues of $\frac{1}{n}X^\top X$ translates into an ellipsoid constraint on the mean parameter sequence [16, 17]. For this reason, Theorem 3 may be viewed as “regression version” of ℓ_2 error bounds for the sequence model under an ellipsoid constraint (cf. Pinsker’s Theorem, [16, 17]). Due to the fact that the latter problem has a very well developed literature, there may be various “neighboring results” elsewhere. Nevertheless, we could not find a direct reference for our stated MSPE bound in the current setup. For the purposes of our work in this paper, the more important point to take away from Theorem 3 is that it can be coupled with Theorem 2 for proving consistency of RB.

3.4 Confidence intervals for mean responses, conditionally on a Gaussian design

In this section, we consider the situation where the design matrix X has rows $X_i^\top \in \mathbb{R}^p$ drawn i.i.d. from a multivariate normal distribution $N(0, \Sigma)$, with $X \perp\!\!\!\perp \varepsilon$. (The covariance matrix Σ may vary with n .) Conditionally on a realization of X , we analyze the RB approximation of the laws $\Psi_\rho(F_0; X_i) = \mathcal{L}(X_i^\top(\widehat{\beta}_\rho - \beta)|X)$. As discussed in Section 1, this corresponds to the problem of setting confidence intervals for the mean responses $\mathbb{E}[Y_i] = X_i^\top \beta$. Assuming that the population eigenvalues $\lambda_i(\Sigma)$ obey a decay condition, we show below in Theorem 4 that RB succeeds with high \mathbb{P}_X -probability. Moreover, this consistency statement holds for all of the laws $\Psi_\rho(F_0; X_i)$ *simultaneously*. That is, among the n distinct laws $\Psi_\rho(F_0; X_i)$, $i = 1, \dots, n$, even the worst bootstrap approximation is still consistent. We now state some population-level assumptions.

A4. The operator norm of $\Sigma \in \mathbb{R}^{p \times p}$ satisfies $\|\Sigma\|_{\text{op}} \lesssim 1$.

Next, we impose a decay condition on the eigenvalues of Σ . This condition also ensures that Σ is invertible for each fixed p — even though the bottom eigenvalue may become arbitrarily small as p becomes large. It is important to notice that we now use η for the decay exponent of the population eigenvalues, whereas we used ν when describing the sample eigenvalues in the previous section.

A5. There is a number $\eta > 0$, and absolute constants $k_1, k_2 > 0$, such that for all $i = 1, \dots, p$,

$$k_1 i^{-\eta} \leq \lambda_i(\Sigma) \leq k_2 i^{-\eta}.$$

A6. There are absolute constants $k_3, k_4 \in (0, 1)$ such that for all $n \geq 3$, we have the bounds $k_3 \leq \frac{p}{n} \leq k_4$ and $p \leq n - 2$.

The following lemma collects most of the effort needed in proving our final result in Theorem 4. Here it is also helpful to recall the notation $\rho/n = n^{-\gamma}$ and $\varrho/n = n^{-\theta}$ from Assumption 2.

Lemma 2. Suppose that the matrix $X \in \mathbb{R}^{n \times p}$ has rows X_i^\top drawn i.i.d. from $N(0, \Sigma)$, and that Assumptions 2–6 hold. Furthermore, assume that γ chosen so that $0 < \gamma < \min\{\eta, 1\}$. Then, the statements below are true.

(i) (bias inequality)

Fix any $\tau > 0$. Then, there is an absolute constant $\kappa_0 > 0$, such that for all large n , the following event holds with \mathbb{P}_X -probability at least $1 - n^{-\tau} - ne^{-n/16}$,

$$\max_{1 \leq i \leq n} b_\rho^2(X; X_i) \leq \kappa_0 \cdot n^{-\gamma} \cdot (\tau + 1) \log(n + 2). \quad (15)$$

(ii) (variance inequality)

There are absolute constants $\kappa_1, \kappa_2 > 0$ such that for all large n , the following event holds with \mathbb{P}_X -probability at least $1 - 4n \exp(-\kappa_1 n^{\frac{2}{\eta}})$,

$$\max_{1 \leq i \leq n} \frac{1}{v_\rho(X; X_i)} \leq \kappa_2 n^{1 - \frac{2}{\eta}}. \quad (16)$$

(iii) (mspe inequalities)

Suppose that θ is chosen as $\theta = 2\eta/3$ when $\eta \in (0, \frac{1}{2})$, and that θ is chosen as $\theta = \frac{\eta}{1+\eta}$ when $\eta > \frac{1}{2}$. Then, there are absolute constants $\kappa_3, \kappa_4, \kappa_5, \kappa_6 > 0$ such that for all large n ,

$$\text{mspe}(\hat{\beta}_\varrho | X) \leq \begin{cases} \kappa_4 n^{-\frac{2\eta}{3}} & \text{with } \mathbb{P}_X\text{-probability at least } 1 - \exp(-\kappa_3 n^{2-4\eta/3}), \text{ if } \eta \in (0, \frac{1}{2}) \\ \kappa_6 n^{-\frac{\eta}{\eta+1}} & \text{with } \mathbb{P}_X\text{-probability at least } 1 - \exp(-\kappa_5 n^{\frac{2}{1+\eta}}), \text{ if } \eta > \frac{1}{2}. \end{cases}$$

Remarks. Note that the two rates in part (iii) coincide as η approaches $1/2$. At a conceptual level, the entire lemma may be explained in relatively simple terms. Viewing the quantities $\text{mspe}(\hat{\beta}_\varrho | X)$, $b_\rho^2(X; X_i)$ and $v_\rho(X; X_i)$ as functionals of a Gaussian matrix, the proof involves deriving concentration bounds for each of them. Indeed, this is plausible given that these quantities are smooth functionals of X . However, the difficulty of the proof arises from the fact that they are also highly non-linear functionals of X .

We now combine Lemmas 1 and 2 with Theorems 1 and 2 to show that all of the laws $\Psi_\rho(F_0; X_i)$ can be simultaneously approximated via our two-stage RB method.

Theorem 4. Suppose that F_0 has a finite fourth moment, Assumptions 2–6 hold, and γ is chosen so that $\frac{\eta}{1+\eta} < \gamma < \min\{\eta, 1\}$. Also suppose that θ is chosen as $\theta = 2\eta/3$ when $\eta \in (0, \frac{1}{2})$, and $\theta = \frac{\eta}{\eta+1}$ when $\eta > \frac{1}{2}$. Then, there is a sequence of positive numbers δ_n with $\lim_{n \rightarrow \infty} \delta_n = 0$, such that the event

$$\mathbb{E} \left[\max_{1 \leq i \leq n} d_2^2 \left(\frac{1}{\sqrt{v_\rho}} \Psi_\rho(F_0; X_i), \frac{1}{\sqrt{v_\rho}} \Phi_\rho(\hat{F}_\varrho; X_i) \right) \mid X \right] \leq \delta_n \quad (17)$$

has \mathbb{P}_X -probability tending to 1 as $n \rightarrow \infty$.

Remark. Lemma 2 gives explicit bounds on the numbers δ_n , as well as the probabilities of the corresponding events, but we have stated the result in this way for the sake of readability.

4 Simulations

In four different settings of n, p , and the decay parameter η , we compared the nominal 90% confidence intervals (CIs) of four methods: “oracle”, “ridge”, “normal”, and “OLS”, to be described below. In each setting, we generated $N_1 := 100$ random designs X with i.i.d. rows drawn from $N(0, \Sigma)$, where $\lambda_j(\Sigma) = j^{-\eta}$, $j = 1, \dots, p$, and the eigenvectors of Σ were drawn randomly by setting them to be the Q factor in a QR decomposition of a standard $p \times p$ Gaussian matrix. Then, for each realization of X , we generated $N_2 := 1000$ realizations of Y according to the model (1), where $\beta = \mathbf{1}/\|\mathbf{1}\|_2 \in \mathbb{R}^p$, and F_0 is the centered t distribution on 5 degrees of freedom, rescaled to have standard deviation $\sigma = 0.1$. For each X , and each corresponding Y , we considered the problem of setting a 90% CI for the mean response value $X_{i^*}^\top \beta$, where X_{i^*} is the row with the highest leverage score, i.e. $i^* = \operatorname{argmax}_{1 \leq i \leq n} H_{ii}$ and $H := X(X^\top X)^{-1}X^\top$. This problem was shown in B&F 1983 to be a case where the standard RB method based on least-squares fails when $p/n \asymp 1$. Below, we refer to this method as “OLS”.

To describe the other three methods, “ridge” refers to the interval $[X_{i^*}^\top \hat{\beta}_\rho - \hat{q}_{0.95}, X_{i^*}^\top \hat{\beta}_\rho - \hat{q}_{0.05}]$, where \hat{q}_α is the $\alpha\%$ quantile of the numbers z_1, \dots, z_B computed in the proposed algorithm in Section 2, with $B = 1000$ and $c^\top = X_{i^*}^\top$. To choose the parameters ρ and ϱ for a given X and Y , we first computed \hat{r} as the value that optimized the MSPE error of a ridge estimator $\hat{\beta}_r$ with respect to 5-fold cross validation; i.e. cross validation was performed for every distinct pair (X, Y) . We then put $\varrho = 5\hat{r}$ and $\rho = 0.1\hat{r}$, as we found the prefactors 5 and 0.1 to work adequately across various settings. (Optimizing ϱ with respect to MSPE is motivated by Theorems 1, 2, and 3. Also, choosing ρ to be somewhat smaller than ϱ conforms with the constraints on θ and γ in Theorem 4.) The method “normal” refers to the CI based on the normal approximation $\mathcal{L}(X_{i^*}^\top(\hat{\beta}_\rho - \beta)|X) \approx N(0, \hat{\tau}^2)$, where $\hat{\tau}^2 = \hat{\sigma}^2 \|X_{i^*}^\top (X^\top X + \rho I_{p \times p})^{-1} X^\top\|_2^2$, $\rho = 0.1\hat{r}$, and $\hat{\sigma}^2$ is the usual unbiased estimate of σ^2 based on OLS residuals. The “oracle” method refers to the interval $[X_{i^*}^\top \hat{\beta}_\rho - \tilde{q}_{0.95}, X_{i^*}^\top \hat{\beta}_\rho - \tilde{q}_{0.05}]$, with $\rho = 0.1\hat{r}$, and \tilde{q}_α being the empirical $\alpha\%$ quantile of $X_{i^*}^\top(\hat{\beta}_\rho - \beta)$ over all 1000 realizations of Y based on a given X . (This accounts for the randomness in $\rho = 0.1\hat{r}$.)

Within a given setting of the triplet (n, p, η) , we refer to the “coverage” of a method as the fraction of the $N_1 \times N_2 = 10^5$ instances where the method’s CI contained the parameter $X_{i^*}^\top \beta$. Also, we refer to “width” as the average width of a method’s intervals over all of the 10^5 instances. The four settings of (n, p, η) correspond to moderate/high dimension and moderate/fast decay of the eigenvalues $\lambda_i(\Sigma)$. Even in the moderate case of $p/n = 0.45$, the results show that the OLS intervals are too narrow and have coverage noticeably less than 90%. As expected, this effect becomes more pronounced when $p/n = 0.95$. The ridge and normal intervals perform reasonably well across settings, with both performing much better than OLS. However, it should be emphasized that our study of RB is motivated by the desire to gain insight into the behavior of the bootstrap in high dimensions — rather than trying to outperform particular methods. In future work, we plan to investigate the relative merits of the ridge and normal intervals in greater detail.

Table 1: Comparison of nominal 90% confidence intervals

		oracle	ridge	normal	OLS
setting 1 $n = 100, p = 45, \eta = 0.5$	width	0.21	0.20	0.23	0.16
	coverage	0.90	0.87	0.91	0.81
setting 2 $n = 100, p = 95, \eta = 0.5$	width	0.22	0.26	0.26	0.06
	coverage	0.90	0.88	0.88	0.42
setting 3 $n = 100, p = 45, \eta = 1$	width	0.20	0.21	0.22	0.16
	coverage	0.90	0.90	0.91	0.81
setting 4 $n = 100, p = 95, \eta = 1$	width	0.21	0.26	0.23	0.06
	coverage	0.90	0.92	0.87	0.42

Acknowledgements. MEL thanks Prof. Peter J. Bickel for many helpful discussions, and gratefully acknowledges the DOE CSGF under grant DE-FG02-97ER25308, as well as the NSF-GRFP.

A Proof of Theorem 1

Proof. Due to line (4) and Lemma 8.8 in B&F 1981,

$$d_2^2(\Psi_\rho(F_0; c), \Phi_\rho(\hat{F}; c)) = d_2^2(\Phi_\rho(F_0; c), \Phi_\rho(\hat{F}; c)) + (c^\top \delta(X))^2. \quad (18)$$

If $\varepsilon^* \in \mathbb{R}^n$ is a random vector whose entries are drawn i.i.d. from \hat{F} , then the definition of Φ_ρ gives the matching relations

$$\begin{aligned} \Phi_\rho(F_0; c) &= \mathcal{L}(c^\top (X^\top X + \rho)^{-1} X^\top \varepsilon \mid X) \\ \Phi_\rho(\hat{F}; c) &= \mathcal{L}(c^\top (X^\top X + \rho)^{-1} X^\top \varepsilon^* \mid X). \end{aligned} \quad (19)$$

To make use of these relations, we apply Lemma 8.9 in B&F 1981, which implies that if $w \in \mathbb{R}^n$ is a generic deterministic vector, and if $U = (U_1, \dots, U_n)$ and $V = (V_1, \dots, V_n)$ are random vectors with i.i.d. entries, then

$$d_2^2(w^\top U, w^\top V) \leq \|w\|_2^2 \cdot d_2(U_1, V_1)^2.$$

Therefore,

$$\begin{aligned} d_2^2(\Phi_\rho(F_0), \Phi_\rho(\hat{F})) &\leq \|c^\top (X^\top X + \rho)^{-1} X^\top\|_2^2 \cdot d_2^2(\varepsilon_1, \varepsilon_1^*) \\ &= \frac{1}{\sigma^2} \cdot v_\rho(X; c) \cdot d_2^2(F_0, \hat{F}). \end{aligned} \quad (20)$$

Combining this with line (18) and dividing through by v_ρ proves the claim. \square

B Proof of Theorem 2

Proof. By the triangle inequality,

$$d_2^2(\hat{F}, F_0) \leq 2 d_2^2(\hat{F}, F_n) + 2 d_2^2(F_n, F_0). \quad (21)$$

Let \tilde{F}_n be the (uncentered) empirical distribution of the residuals \hat{e} of $\hat{\beta}$, which places mass $1/n$ at each value \hat{e}_i , for $i = 1, \dots, n$. The proofs of Lemmas 2.1 and 2.2 in Freedman 1981, show that

$$\begin{aligned} \mathbb{E}[d_2^2(\hat{F}, F_n) \mid X] &\leq \mathbb{E}\left[\left(\frac{1}{n} \sum_{i=1}^n \varepsilon_i\right)^2\right] + \mathbb{E}[d_2^2(\tilde{F}_n, F_n) \mid X] \\ &\leq \frac{1}{n} \sigma^2 + \frac{1}{n} \mathbb{E}[\|\hat{e} - \varepsilon\|_2^2 \mid X] \\ &= \frac{1}{n} \sigma^2 + \frac{1}{n} \mathbb{E}[\|X(\beta - \hat{\beta})\|_2^2 \mid X], \end{aligned} \quad (22)$$

where we have used the algebraic identity $\hat{e} - \varepsilon = X(\beta - \hat{\beta})$, which holds for any estimator $\hat{\beta}$. This completes the proof. \square

C Proof of Theorem 3

Proof. We begin with a simple bias-variance decomposition,

$$\text{mspe}(\hat{\beta}_\varrho | X) = \frac{1}{n} \mathbb{E}\left[\|X(\hat{\beta}_\varrho - \mathbb{E}[\hat{\beta}_\varrho | X])\|_2^2 \mid X\right] + \frac{1}{n} \|X(\mathbb{E}[\hat{\beta}_\varrho | X] - \beta)\|_2^2. \quad (23)$$

We will handle the bias and variance terms separately. To consider the bias term, note that $\mathbb{E}[\hat{\beta}_\rho | X] - \beta = B\beta$, where

$$B = (X^\top X + \varrho I_{p \times p})^{-1} X^\top X - I_{p \times p}.$$

Hence,

$$\begin{aligned} \frac{1}{n} \|X(\mathbb{E}[\hat{\beta}_\varrho | X] - \beta)\|_2^2 &= \frac{1}{n} \|XB\beta\|_2^2 \\ &= \beta^\top B \left(\frac{1}{n} X^\top X\right) B\beta. \end{aligned} \quad (24)$$

If we let $l_i = \lambda_i(\frac{1}{n}X^\top X)$, then the eigenvalues of $B(\frac{1}{n}X^\top X)B$ are of the form $\mu_i := \frac{l_i(\varrho/n)^2}{(l_i + \varrho/n)^2}$. In particular, it is simple to check¹ that $\max_i \mu_i \asymp \varrho/n$ whenever $\theta \leq \nu$, and so

$$\beta^\top B(\frac{1}{n}X^\top X)B\beta \lesssim \frac{\varrho}{n}\|\beta\|_2^2 = n^{-\theta}\|\beta\|_2^2. \quad (25)$$

Note that this bound is tight, since it is achieved whenever β is parallel to the top eigenvector of $B(\frac{1}{n}X^\top X)B$.

To consider the variance term, note that $\widehat{\beta}_\varrho - \mathbb{E}[\widehat{\beta}_\varrho|X] = (X^\top X + \varrho I_{p \times p})^{-1}X^\top \varepsilon$, and so

$$\begin{aligned} \frac{1}{n}\mathbb{E}\left[\|X(\widehat{\beta}_\varrho - \mathbb{E}[\widehat{\beta}_\varrho|X])\|_2^2 \mid X\right] &= \frac{1}{n}\text{tr}\left(\left(X^\top X(X^\top X + \varrho I_{p \times p})^{-1}\right)^2\right) \\ &= \frac{1}{n}\sum_{i=1}^{n \wedge p} \left(\frac{l_i}{l_i + \varrho/n}\right)^2. \end{aligned} \quad (26)$$

It is natural to decompose the sum in terms of the index set

$$\mathcal{I}(n) := \{i \in \{1, \dots, n \wedge p\} : l_i \geq \varrho/n\}, \quad (27)$$

which satisfies $|\mathcal{I}(n)| \asymp n^{\frac{\theta}{\nu}}$. We will bound the variance term in two complementary cases; either $\nu > 1/2$ or $\nu \leq 1/2$. First assume $\nu > 1/2$. Then,

$$\frac{1}{n}\sum_{i=1}^p \left(\frac{l_i}{l_i + \varrho/n}\right)^2 = \frac{1}{n}\sum_{i \in \mathcal{I}(n)} \left(\frac{l_i}{l_i + \varrho/n}\right)^2 + \frac{1}{n}\sum_{i \notin \mathcal{I}(n)} \left(\frac{l_i}{l_i + \varrho/n}\right)^2 \quad (28)$$

$$\lesssim \frac{1}{n}|\mathcal{I}(n)| + \frac{1}{n}\int_{|\mathcal{I}(n)|}^{n \wedge p} \frac{x^{-2\nu}}{(\varrho/n)^2} dx \quad (29)$$

$$\lesssim n^{-1}\left(n^{\frac{\theta}{\nu}} + n^{2\theta} \cdot (|\mathcal{I}(n)|)^{1-2\nu}\right) \quad \text{using } \nu > \frac{1}{2} \quad (30)$$

$$\asymp n^{-1}\left(n^{\frac{\theta}{\nu}} + n^{2\theta} \cdot (n^{\frac{\theta}{\nu}})^{(1-2\nu)}\right) \quad (31)$$

$$= 2n^{\frac{\theta-\nu}{\nu}}. \quad (32)$$

To see that this upper bound is tight, note that in line (28), we can use the term-wise lower bounds

$$\left(\frac{l_i}{l_i + \varrho/n}\right)^2 \geq \begin{cases} \frac{1}{4} & \text{if } i \in \mathcal{I}(n) \\ \frac{1}{4} \frac{l_i^2}{(\varrho/n)^2} & \text{if } i \notin \mathcal{I}(n), \end{cases} \quad (33)$$

and then apply an integral approximation from below (which leads to the same rate). Combining the bias and variance pieces, we have shown that

$$\frac{1}{n}\mathbb{E}\|X(\widehat{\beta}_\varrho - \beta)\|_2^2 \lesssim n^{\frac{\theta-\nu}{\nu}} + n^{-\theta} \quad \text{if } \nu > \frac{1}{2}.$$

The bound is optimized when the two terms on the right side have the same rate, which leads to the choice $\theta = \frac{\nu}{\nu+1}$.

In the case where $\nu \in (0, \frac{1}{2})$, the calculation proceeds in the same way up to line (30), where we obtain the bound

$$\frac{1}{n}\sum_{i=1}^{n \wedge p} \left(\frac{l_i}{l_i + \varrho/n}\right)^2 \lesssim n^{-1}\left(n^{\frac{\theta}{\nu}} + n^{2\theta} \cdot n^{1-2\nu}\right) \quad (34)$$

$$= n^{\frac{\theta-\nu}{\nu}} + n^{2(\theta-\nu)}. \quad (35)$$

This bound is also tight due to the same reasoning as above. Note that in order for the bound to tend to 0 as $n \rightarrow \infty$, we must choose $\theta < \nu$. Furthermore, since we are working under the assumption $\nu \in (0, \frac{1}{2})$, it follows that the right side of line (35) has rate equal to $n^{2(\theta-\nu)}$. Combining the rates for the bias and variance shows that

$$\frac{1}{n}\mathbb{E}\|X(\widehat{\beta}_\varrho - \beta)\|_2^2 \lesssim n^{2(\theta-\nu)} + n^{-\theta} \quad \text{if } \nu \in (0, \frac{1}{2}).$$

The bound is optimized when the two terms on the right side have the same rate, which leads to the choice $\theta = \frac{2\nu}{3}$. \square

¹Note that if $t \in \mathbb{R}$ and $f(t) := \frac{t(\varrho/n)^2}{(t + \varrho/n)^2}$, then f is maximized at $t = \varrho/n$. Also, if $\theta \leq \nu$, then there at least one l_i that scales at the rate of ϱ/n .

D Proof of Lemma 2

The proof is split up into three pieces, corresponding to parts (i), (ii), and (iii) in the statement of the result.

D.1 The bias inequality (15)

We prove inequality (15) by combining Lemmas 3 and 5 below.

Lemma 3. *Assume the conditions of Lemma 2. For each $i \in \{1, \dots, n\}$, there are independent random vectors $u_i(X), w(X) \in \mathbb{R}^p$ such that the random variable $X_i^\top \delta(X)$ can be represented algebraically as*

$$b_\rho(X; X_i) = X_i^\top \delta(X) = u_i(X)^\top w(X).$$

Here, the vectors $u_i(X)$ can be represented in law as

$$u_i(X) \stackrel{\mathcal{L}}{=} \frac{1}{\|z\|_2} \Pi_p(z), \quad (36)$$

where $z \in \mathbb{R}^n$ is a standard Gaussian vector, and $\Pi_p(z) := (z_1, \dots, z_p)$. Also, the vector $w(X)$ satisfies the bound $\|w(X)\|_2^2 \leq \frac{\rho}{4} \|\beta\|_2^2$ almost surely.

Proof. To fix notation, we write $X^\top = \Sigma^{1/2} Z^\top$ where $Z^\top \in \mathbb{R}^{p \times n}$ is a standard Gaussian matrix. Recall that $\delta(X) = B\beta$, where

$$B = I_{p \times p} - (X^\top X + \rho I_{p \times p})^{-1} X^\top X.$$

Let $Z = HLG^\top$ be a “signed s.v.d.” for Z , as defined in Appendix E.2, where $H \in \mathbb{R}^{n \times p}$, $L \in \mathbb{R}^{p \times p}$, and $G \in \mathbb{R}^{p \times p}$. Now define $u_i(X)$ and $w(X)$ according to

$$X_i^\top \delta(X) = e_i^\top X B \beta = e_i^\top Z \Sigma^{1/2} B \beta = \underbrace{e_i^\top H}_{=: u_i(X)^\top} \underbrace{L G^\top \Sigma^{1/2} B \beta}_{=: w(X)}. \quad (37)$$

From Lemma 17 in Appendix E.2, the rows $e_i^\top H$ can be represented in distribution as $\frac{1}{\|z\|_2} \Pi_p(z)$. The same lemma also shows that the three matrices H , L , and G are independent.

Hence, to show that $u_i(X)$ and $w(X)$ are independent, it suffices to show that $w(X)$ is a function only of G and L . In turn, it is enough to show that B is a function only of G and L . But this is simple, because B is a function only of the matrix $X^\top X$, which may be written as

$$X^\top X = \Sigma^{1/2} Z^\top Z \Sigma^{1/2} = \Sigma^{1/2} G L^2 G^\top \Sigma^{1/2}. \quad (38)$$

It remains to show that $\|w(X)\|_2^2 \leq \frac{\rho}{4} \|\beta\|_2^2$ almost surely. Combining the definition of $w(X)$ with line (38) gives

$$\|w(X)\|_2^2 = \beta^\top (B X^\top X B) \beta. \quad (39)$$

The eigenvalues of $B X^\top X B$ are of the form $\mu_i := n \frac{(\rho/n)^2 l_i}{(l_i + \rho/n)^2}$ where $l_i = \lambda_i(\frac{1}{n} X^\top X)$, and it is simple to check that the inequality $\max_i \mu_i \leq \frac{\rho}{4}$ holds for every realization of X . \square

Before proceeding to the second portion of the proof of inequality (15), we record some well-known tail bounds for Gaussian quadratic forms due to Laurent and Massart [18], which will be useful at various points later on.

Lemma 4 (Laurent & Massart, 2001). *Let $A \in \mathbb{R}^{n \times n}$ be a fixed symmetric matrix, and let $z \in \mathbb{R}^n$ be a standard Gaussian vector. Then, for every $t > 0$,*

$$\mathbb{P}\left[z^\top A z \geq \text{tr}(A) + 2 \|A\|_F \sqrt{t} + 2 \|A\|_{\text{op}} t\right] \leq \exp(-t) \quad (40)$$

and

$$\mathbb{P}\left[z^\top A z \leq \text{tr}(A) - 2 \|A\|_F \sqrt{t}\right] \leq \exp(-t). \quad (41)$$

The next lemma completes the proof of inequality (15).

Lemma 5. *Assume the conditions of Lemma 2, and let $\tau > 0$ be a constant. Then for every $n \geq 1$, the following event holds with probability at least $1 - n^{-\tau} - ne^{-n/16}$,*

$$\max_{1 \leq i \leq n} b_\rho^2(X; X_i) \leq 5\|\beta\|_2^2 \cdot n^{-\gamma} \cdot (\tau + 1) \log(n + 2). \quad (42)$$

Proof. Applying the representation for $b_\rho(X; X_i)$ given in Lemma 3, there is a standard Gaussian vector $z \in \mathbb{R}^n$, such that $u_i(X) \stackrel{\mathcal{L}}{=} \Pi_p(z)/\|z\|_2$. Consequently,

$$b_\rho^2(X; X_i) \stackrel{\mathcal{L}}{=} \frac{1}{\|z\|_2^2} \cdot \Pi_p(z) \left(w(X) w(X)^\top \right) \Pi_p(z), \quad (43)$$

where we may take z and $w(X)$ to be independent by the same lemma. Using Lemma 4 on Gaussian quadratic forms, as well as the fact that $\|w(X)\|_2^2 \leq \frac{\rho}{4}\|\beta\|_2^2$ almost surely, we have for all $t > 0$,

$$\mathbb{P} \left[\Pi_p(z)^\top \left(w(X) w(X)^\top \right) \Pi_p(z) \geq \frac{\rho}{4}\|\beta\|_2^2 (1 + 2\sqrt{t} + 2t) \mid w(X) \right] \leq \exp(-t). \quad (44)$$

The same lemma also implies that for all $t' \in (0, \frac{1}{4})$,

$$\mathbb{P} \left[\frac{1}{\|z\|_2^2} \geq \frac{1}{(1 - 2\sqrt{t'})n} \right] \leq \exp(-nt'). \quad (45)$$

Now, we combine the bounds by integrating out $w(X)$ in line (44) and choosing $t' = 1/16$ in line (45). Taking a union bound, we conclude that for any $t > 0$, and any fixed $i = 1, \dots, n$,

$$\mathbb{P} \left[b_\rho^2(X; X_i) \leq \frac{\rho}{n} \cdot \|\beta\|_2^2 \cdot \frac{1}{2} (1 + 2\sqrt{t} + 2t) \right] \geq 1 - e^{-t} - e^{-n/16}. \quad (46)$$

Finally, another union bound shows that the maximum of the $b_\rho(X, X_i^\top)$ satisfies

$$\mathbb{P} \left[\max_{1 \leq i \leq n} b_\rho^2(X; X_i) \leq \frac{\rho}{n} \cdot \|\beta\|_2^2 \cdot \frac{1}{2} (1 + 2\sqrt{t} + 2t) \right] \geq 1 - e^{-t + \log(n)} - ne^{-n/16}, \quad (47)$$

which implies the stated result after choosing $t = (\tau + 1) \log(n + 2)$, and noting that since $t \geq 1$, we have $\frac{1}{2}(1 + 2\sqrt{t} + 2t) \leq 5t = 5(\tau + 1) \log(n + 2)$, as well as $e^{-t + \log(n)} \leq e^{-\tau \log(n + 2)} \leq n^{-\tau}$ for every $n \geq 1$. \square

D.2 The variance inequality (16)

The following “representation lemma” will serve as the basis for controlling the variance $v_\rho(X; X_i) = \sigma^2 \|X_i^\top (X^\top X + \rho I_{p \times p})^{-1} X^\top\|_2^2$.

Lemma 6. *Assume the conditions of Lemma 2. For each $i \in \{1, \dots, n\}$, there is a random vector $v_i(X) \in \mathbb{R}^p$ and a random matrix $M(X) \in \mathbb{R}^{p \times p}$ that are independent and satisfy the algebraic relation*

$$\|X_i^\top (X^\top X + \rho I_{p \times p})^{-1} X^\top\|_2^2 = v_i(X)^\top M(X) v_i(X).$$

Here, the vector $v_i(X)$ can be represented in law as

$$v_i(X) \stackrel{\mathcal{L}}{=} \frac{1}{\|z\|_2} \Pi_p(z), \quad (48)$$

where $z \in \mathbb{R}^n$ is a standard Gaussian vector and $\Pi_p(z) = (z_1, \dots, z_p)$. Also, the matrix $M(X)$ satisfies the algebraic relation

$$\text{tr}(M(X)) = \|X(X^\top X + \rho I_{p \times p})^{-1} X^\top\|_F^2. \quad (49)$$

An explicit formula for $M(X)$ is given below.

Proof. Define the matrix $A := (X^\top X + \rho I_{p \times p})^{-1} X^\top X (X^\top X + \rho I_{p \times p})^{-1}$. Then,

$$\|X_i^\top (X^\top X + \rho I_{p \times p})^{-1} X^\top\|_2^2 = e_i^\top X A X^\top e_i. \quad (50)$$

Using the notation in the proof of the previous lemma, let $X = Z\Sigma^{1/2}$ where $Z \in \mathbb{R}^{n \times p}$ is a standard Gaussian random matrix. Furthermore, let $Z = HLG^\top$ be a signed s.v.d. for Z , as defined in Appendix E.2. Then, we define $v_i(X)$ and $M(X)$ according to

$$e_i^\top X A X^\top e_i = \underbrace{e_i^\top H}_{=: v_i(X)^\top} \underbrace{L G^\top \Sigma^{1/2} A \Sigma^{1/2} G L^\top}_{=: M(X)} H^\top e_i. \quad (51)$$

Some algebra shows that $M(X)$ satisfies the relation (49). As in the proof of Lemma 3, the argument is completed using two properties of the signed s.v.d. of a standard Gaussian matrix: The rows of H can be represented as $\Pi_p(z)/\|z\|_2$ where $z \in \mathbb{R}^n$ is a standard Gaussian vector, and the matrices H , L , and G^\top are independent. (See Lemma 17 in Appendix E.2.) To show that $v_i(X)$ and $M(X)$ are independent, first note that $v_i(X)$ only depends on H . Also, it is simple to check that $M(X)$ only depends on G and L , because A is a function only of $X^\top X = \Sigma^{1/2} G L^2 G^\top \Sigma^{1/2}$. \square

D.2.1 Concentration of the variance and bounds on its expected value

Due to Lemma 6, for each $i = 1, \dots, n$, we have the representation

$$v_\rho(X; X_i) \stackrel{c}{=} \frac{1}{\|z\|_2^2} \Pi_p(z)^\top M(X) \Pi_p(z), \quad (52)$$

where $z \in \mathbb{R}^n$ is a standard Gaussian vector, independent of $M(X)$. Conditionally on $M(X)$, the quadratic form $\Pi_p(z)^\top M(X) \Pi_p(z)$ concentrates around $\text{tr}(M(X))$ by Lemma 4. The same lemma also implies that $\|z\|_2^2$ concentrates around n . In the next three subsections, we will show that $\sqrt{\text{tr}(M(X))}$ concentrates around its expected value, and obtain upper and lower bounds on the expected value. We will need two-sided bounds in preparation for Theorem 4.

D.2.2 Concentration of $\sqrt{\text{tr}(M(X))}$

Lemma 7. Assume the conditions of Lemma 2. Then for every $t > 0$, and every $n \geq 1$,

$$\mathbb{P} \left[\left| \sqrt{\text{tr}(M(X))} - \mathbb{E} \sqrt{\text{tr}(M(X))} \right| \geq t \right] \leq 2 \exp(-\frac{64}{54} \frac{n^{1-\gamma} t^2}{\|\Sigma\|_{\text{op}}}). \quad (53)$$

Proof. We will show that $\sqrt{\text{tr}(M(X))}$ is a Lipschitz function of a standard Gaussian matrix. Define the function $g_\rho : \mathbb{R}_+ \rightarrow [0, 1]$ by $g_\rho(s) = \frac{s^2}{s^2 + (\rho/n)}$, which satisfies the Lipschitz condition

$$|g_\rho(s) - g_\rho(s')| \leq \mathfrak{L}_n |s - s'|,$$

for all $s, s' \geq 0$, where $\mathfrak{L}_n := \frac{3\sqrt{3}}{8} \frac{1}{\sqrt{\rho/n}}$.

If $\sigma(A) = (\sigma_1(A), \dots, \sigma_k(A))$ denotes the vector of singular values of a rank k matrix A , then we define g_ρ to act on $\sigma(A)$ component-wise, i.e. $g_\rho(\sigma(A)) = (g_\rho(\sigma_1(A)), \dots, g_\rho(\sigma_k(A)))$. Recall from Lemma 6 that

$$\sqrt{\text{tr}(M(X))} = \|X(X^\top X + \rho I_{p \times p})^{-1} X^\top\|_F \quad (54)$$

and note that the i th singular value of the matrix $X(X^\top X + \rho I_{p \times p})^{-1} X^\top$ is given by $g_\rho(\sigma_i(\frac{1}{\sqrt{n}}X))$. Viewing the Frobenius norm of a matrix as the ℓ_2 norm of its singular values, we have

$$\sqrt{\text{tr}(M(X))} = \|g_\rho(\sigma(\frac{1}{\sqrt{n}}X))\|_2. \quad (55)$$

Write $X^\top = \Sigma^{1/2} Z^\top$ for a standard Gaussian matrix $Z \in \mathbb{R}^{n \times p}$, and let $f : \mathbb{R}^{n \times p} \rightarrow \mathbb{R}$ be defined according to

$$f(Z) := \sqrt{\text{tr}(M(X))}.$$

We claim that f is Lipschitz with respect to the Frobenius norm. Let $W^\top \in \mathbb{R}^{p \times n}$ be a generic matrix, and put $A = \frac{1}{\sqrt{n}} \Sigma^{1/2} Z^\top$ and $B = \frac{1}{\sqrt{n}} \Sigma^{1/2} W^\top$. Then,

$$|f(Z) - f(W)| = \left| \|g_\rho(\sigma(A))\|_2 - \|g_\rho(\sigma(B))\|_2 \right| \quad (56)$$

$$\leq \|g_\rho(\sigma(A)) - g_\rho(\sigma(B))\|_2 \quad (57)$$

$$\leq \mathfrak{L}_n \|\sigma(A) - \sigma(B)\|_2 \quad (58)$$

$$\leq \mathfrak{L}_n \|A - B\|_F \quad (\text{Weilandt-Hoffman}) \quad (59)$$

$$= \mathfrak{L}_n \left\| \frac{1}{\sqrt{n}} \Sigma^{1/2} (Z^\top - W^\top) \right\|_F \quad (60)$$

$$\leq \frac{\mathfrak{L}_n}{\sqrt{n}} \sqrt{\|\Sigma\|_{\text{op}}} \cdot \|Z^\top - W^\top\|_F, \quad (61)$$

where we have used a version of the Weilandt-Hoffman inequality for singular values [19, p.186], as well as the inequality $\|M_1 M_2\|_F \leq \|M_1\|_{\text{op}} \|M_2\|_F$, which holds for any square matrix M_1 that is compatible with M_2 . (See Lemma 12 in Appendix E.) The statement of the lemma now follows from the Gaussian concentration inequality. (See Lemma 15 in Appendix E). \square

D.2.3 Upper bound on $\mathbb{E} \sqrt{\text{tr}(M(X))}$

Lemma 8. *Assume the conditions of Lemma 2. Then, the matrix $M(X)$ satisfies*

$$\mathbb{E} \sqrt{\text{tr}(M(X))} \lesssim \begin{cases} n^{(\gamma-\eta)+\frac{1}{2}} & \text{if } \eta \in (0, \frac{1}{2}) \\ n^{\frac{\gamma}{2\eta}} & \text{if } \eta > \frac{1}{2}. \end{cases} \quad (62)$$

Proof. By Jensen's inequality, it is enough to bound $\sqrt{\mathbb{E}[\text{tr}(M(X))]}$ from above. Define the univariate function $\psi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ by $\psi(s) := \frac{s}{(\sqrt{s} + \rho/n)^2}$, and observe that

$$\begin{aligned} \text{tr}(M(X)) &= \text{tr} \left(\left((X^\top X + \rho I_{p \times p})^{-1} X^\top X \right)^2 \right) \\ &= \sum_{i=1}^p \frac{\lambda_i^2(\hat{\Sigma})}{(\lambda_i(\hat{\Sigma}) + \rho/n)^2} \\ &= \sum_{i=1}^p \psi(\lambda_i(\hat{\Sigma}^2)) \\ &= \text{tr}(\psi(\hat{\Sigma}^2)). \end{aligned} \quad (63)$$

Here where we use the “operator calculus” notation $\psi(A) = U\psi(D)U^\top$ where A is a symmetric matrix with spectral decomposition $A = UDU^\top$, and $\psi(D)$ is the diagonal matrix whose i th diagonal entry is $\psi(D_{ii})$. It is simple to check that ψ is a concave, and so $\text{tr}(\psi(\hat{\Sigma}^2))$ is a concave matrix functional of $\hat{\Sigma}^2$ by Lemma 13 in Section E.1 of Appendix A. Therefore, Jensen's inequality implies

$$\begin{aligned} \mathbb{E}[\text{tr}(M(X))] &\leq \text{tr}(\psi(\mathbb{E}[\hat{\Sigma}^2])) \\ &= \sum_{i=1}^p \psi(\lambda_i(\mathfrak{S})), \end{aligned} \quad (64)$$

where we define the matrix $\mathfrak{S} := \mathbb{E}[\hat{\Sigma}^2]$. Since X is Gaussian, $\hat{\Sigma}$ is a Wishart matrix up to scaling, and so Lemma 16 in Appendix E shows that this expectation may be evaluated exactly as

$$\mathfrak{S} = (1 + \frac{1}{n})\Sigma^2 + \frac{\text{tr}(\Sigma)}{n}\Sigma. \quad (65)$$

We will now use this relation to apply an integral approximation to the right side of line (64). Clearly, the eigenvalues of \mathfrak{S} are given by

$$\begin{aligned} \lambda_i(\mathfrak{S}) &= (1 + \frac{1}{n})\lambda_i^2(\Sigma) + \frac{\text{tr}(\Sigma)}{n}\lambda_i(\Sigma) \\ &\asymp i^{-2\eta} + \frac{\text{tr}(\Sigma)}{n}i^{-\eta}. \end{aligned} \quad (66)$$

Let $r \in (0, 1)$ be a constant to be specified later. On the set of indices $1 \leq i \leq \lceil n^r \rceil$ we use the bound $\psi(\lambda_i(\mathfrak{S})) \leq 1$, and on the set of indices $i > \lceil n^r \rceil$ we use the bound $\psi(\lambda_i(\mathfrak{S})) \leq \frac{1}{(\rho/n)^2} \lambda_i(\mathfrak{S})$. Recalling the assumption $\rho/n = n^{-\gamma}$, we may decompose the inequality (64) as²

$$\mathbb{E}[\text{tr}(M(X))] \leq \sum_{i=1}^{\lceil n^r \rceil} \psi(\lambda_i(\mathfrak{S})) + \sum_{i=\lceil n^r \rceil+1}^p \psi(\lambda_i(\mathfrak{S})) \quad (67)$$

$$\lesssim n^r + n^{2\gamma} \int_{n^r}^p (x^{-2\eta} + \frac{\text{tr}(\Sigma)}{n} x^{-\eta}) dx \quad (68)$$

$$=: n^r + n^{2\gamma} h_n(\eta, r). \quad (69)$$

where the function h_n is defined in the last line. The bound is optimized when the two terms on the right are of the same order; i.e. when r solves the rate equation

$$n^r \asymp n^{2\gamma} h_n(\eta, r). \quad (70)$$

Noting that

$$\frac{\text{tr}(\Sigma)}{n} \asymp \begin{cases} n^{-\eta} & \text{if } \eta \in (0, 1) \\ n^{-1} & \text{if } \eta > 1, \end{cases} \quad (71)$$

the quantity $h_n(\eta, r)$ may be computed directly as

$$h_n(\eta, r) \asymp \begin{cases} n^{1-2\eta} & \text{if } \eta \in (0, \frac{1}{2}), \\ n^{r(1-2\eta)} & \text{if } \eta \in (\frac{1}{2}, 1), \\ n^{r(1-2\eta)} + n^{r(1-\eta)-1} & \text{if } \eta > 1. \end{cases} \quad (72)$$

If we let $r = r_*(\eta, \gamma)$ denote the solution of the rate equation (70), then some calculation shows that under the assumption $\gamma \in (0, 1)$,

$$r_*(\eta, \gamma) = \begin{cases} 2(\gamma - \eta) + 1 & \text{if } \eta \in (0, \frac{1}{2}), \\ \frac{\gamma}{\eta} & \text{if } \eta > \frac{1}{2}. \end{cases} \quad (73)$$

When $\eta \in (0, 1)$ this is straightforward. To show the details for $\eta > 1$, note that the rate equation (70) may be written as

$$n^r \asymp n^{2\gamma+r(1-2\eta)} + n^{2\gamma+r(1-\eta)-1}, \quad (74)$$

which is the same as

$$1 \asymp n^{2(\gamma-\eta r)} + n^{2\gamma-\eta r-1}. \quad (75)$$

In order for both terms on the right to be $\mathcal{O}(1)$, the number r must satisfy the constraints

$$r \geq \frac{\gamma}{\eta}, \quad (76)$$

$$r \geq \frac{\gamma}{\eta} + \frac{\gamma-1}{\eta}. \quad (77)$$

Since Lemma 2 assumes $\gamma \in (0, 1)$, only the first constraint matters. Furthermore, when $r \geq \frac{\gamma}{\eta}$, the second term in line (75) is $o(1)$, and we are reduced to choosing r so that $1 \asymp n^{2(\gamma-\eta r)}$, which gives $r = r_*(\eta, \gamma) = \frac{\gamma}{\eta}$. Substituting this value into line (69) completes the proof. (Note from the discussion preceding line (69) that r must lie in the interval $(0, 1)$, and this requires $\gamma/\eta < 1$, which explains the assumption $\gamma < \min\{\eta, 1\}$ in Lemma 2.)

D.2.4 Lower bound on $\mathbb{E}\sqrt{\text{tr}(M(X))}$

Lemma 9. *Assume the conditions of Lemma 2. Then, the matrix $M(X)$ satisfies*

$$\mathbb{E}\sqrt{\text{tr}(M(X))} \gtrsim n^{\frac{\gamma}{2\eta}}. \quad (78)$$

²Note that if $p/n \asymp 1$, it is possible that $n^r > p$ for small values of n . Since we want $n^r \leq p$ for the integral in line (68), Lemma 2 is stated for “all large n ”.

Proof. The variable $\sqrt{\text{tr}(M(X))}$ may be written as $\|X^\top X(X^\top X + \rho I_{p \times p})^{-1}\|_F$. Since the Frobenius norm is a convex matrix functional, Jensen's inequality implies

$$\begin{aligned} \mathbb{E}\sqrt{\text{tr}(M(X))} &\geq \left\| \mathbb{E}\left[X^\top X(X^\top X + \rho I_{p \times p})^{-1}\right] \right\|_F \\ &= \left\| \mathbb{E}\left[(I_{p \times p} + \frac{\rho}{n} \widehat{\Sigma}^{-1})^{-1}\right] \right\|_F, \end{aligned} \quad (79)$$

where the last step follows algebraically with $\widehat{\Sigma} := \frac{1}{n} X^\top X$. If we define the univariate function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ by $f(s) = (1 + \frac{\rho}{n} s)^{-1}$, then last inequality is the same as

$$\mathbb{E}\sqrt{\text{tr}(M(X))} \geq \left\| \mathbb{E}[f(\widehat{\Sigma}^{-1})] \right\|_F. \quad (80)$$

It is a basic fact that f is operator convex on the domain of positive semidefinite matrices [20, p.117]. This yields an operator version of Jensen's inequality with respect to the Loewner ordering (Lemma 14 in Appendix E):

$$\mathbb{E}[f(\widehat{\Sigma}^{-1})] \succeq f(\mathbb{E}[\widehat{\Sigma}^{-1}]). \quad (81)$$

Furthermore, if two matrices satisfy $A \succeq B \succeq 0$, then $\|A\|_F \geq \|B\|_F$ [21, Corollary 7.7.4]. Using this fact, as well as the formula for the expected inverse of a Wishart matrix [22, p. 97], we obtain

$$\begin{aligned} \mathbb{E}\sqrt{\text{tr}(M(X))} &\geq \left\| f(\mathbb{E}[\widehat{\Sigma}^{-1}]) \right\|_F \\ &= \left\| f\left(\frac{n}{n-p-1} \Sigma^{-1}\right) \right\|_F \\ &= \left(\sum_{i=1}^p \frac{1}{\left(1 + \frac{\rho}{n} \cdot \frac{n}{n-p-1} \lambda_i(\Sigma^{-1})\right)^2} \right)^{1/2} \\ &= \left(\sum_{i=1}^p \frac{\lambda_i^2(\Sigma)}{\left(\lambda_i(\Sigma) + \frac{\rho}{n} \cdot \frac{n}{n-p-1}\right)^2} \right)^{1/2}. \end{aligned} \quad (82)$$

Define the index set $J = \{i \in \{1, \dots, p\} : \lambda_i(\Sigma) \geq \frac{\rho}{n} \cdot \frac{n}{n-p-1}\}$. For any $i \in J$, the i th summand in the previous line is at least $1/4$. Also, assumption **A6** that p/n is bounded strictly between 0 and 1, as well as the decay condition on the $\lambda_i(\Sigma)$, imply that $|J| \asymp n^{\gamma/\eta}$, which completes the proof. \square

D.2.5 Putting the variance pieces together

Combining Lemmas 7, 8, and 9 with the Gaussian concentration inequality (Lemma 15 in Section E.2 of Appendix A) immediately gives the following result. (We choose t to be proportional to the relevant bound on $\mathbb{E}[\sqrt{\text{tr}(M(X))}]$ in the Gaussian concentration inequality.)

Lemma 10. *Assume the conditions of Lemma 2 and let $\text{tr}(M(X))$ be as in line (49). Then, there are absolute constants $\kappa_1, \kappa_2, \dots, \kappa_6 > 0$ such that the following upper-tail bounds hold for all large n ,*

$$\mathbb{P}\left[\text{tr}(M(X)) \geq \kappa_1 n^{2(\gamma-\eta)+1}\right] \leq \exp(-\kappa_2 n^{2(1-\eta)+\gamma}), \quad \text{if } \eta \in (0, \frac{1}{2}), \quad (83)$$

and

$$\mathbb{P}\left[\text{tr}(M(X)) \geq \kappa_3 n^{\gamma/\eta}\right] \leq \exp(-\kappa_4 n^{1+\frac{\gamma(1-\eta)}{\eta}}), \quad \text{if } \eta > \frac{1}{2}, \quad (84)$$

and the following lower-tail bound holds for all large n ,

$$\mathbb{P}\left[\text{tr}(M(X)) \leq \kappa_5 n^{\gamma/\eta}\right] \leq \exp(-\kappa_6 n^{1+\frac{\gamma(1-\eta)}{\eta}}), \quad \text{if } \eta > 0. \quad (85)$$

Remarks. Note that in order for the last two probabilities to be small for large values of $\eta > 0$, it is necessary that $\gamma < 1$, as assumed in Lemma 2. The next result completes the assembly of the results in this Subsection D.2. Although the first two bounds in Lemma 11 are not necessary for the statement of Theorem 4, they show that the variance $v_\rho(X; X_i)$ tends 0 as $n \rightarrow \infty$ when $\gamma < \eta$, as assumed in Theorem 4. In other words, we imposed the assumption $\gamma < \eta$ so that confidence intervals based on $\Phi_\rho(\widehat{F}_\rho; X_i)$ have width that tends to 0 asymptotically.

Lemma 11. Assume the conditions of Theorem 4 and let $\text{tr}(M(X))$ be as in line (49). Assume $\gamma < \min\{\eta, 1\}$. Then, there are absolute constants $k_1, k_2, \dots, k_6 > 0$ such that the following upper-tail bounds hold for all large n ,

$$\mathbb{P}\left[\max_{1 \leq i \leq n} v_\rho(X; X_i) \leq k_1 n^{2(\gamma-\eta)}\right] \geq 1 - 4n \exp(-k_2 n^{\frac{\gamma}{\eta}}), \quad \text{if } \eta \in (0, \frac{1}{2}) \quad (86)$$

and

$$\mathbb{P}\left[\max_{1 \leq i \leq n} v_\rho(X; X_i) \leq k_3 n^{\frac{\gamma}{\eta}-1}\right] \geq 1 - 4n \exp(-k_4 n^{\frac{\gamma}{\eta}}), \quad \text{if } \eta > \frac{1}{2}, \quad (87)$$

and

$$\mathbb{P}\left[\max_{1 \leq i \leq n} \frac{1}{v_\rho(X; X_i)} \leq k_5 n^{1-\frac{\gamma}{\eta}}\right] \geq 1 - 4n \exp(-k_6 n^{\frac{\gamma}{\eta}}), \quad \text{if } \eta > 0. \quad (88)$$

Proof. We only prove the last inequality (88), since the other two inequalities are proven in a similar way. By Lemma 6, we have

$$v_\rho(X; X_i) \stackrel{\mathcal{L}}{=} \frac{1}{\|z\|_2^2} \Pi_p(z)^\top M(X) \Pi_p(z) \quad (89)$$

where $z \sim N(0, I_{p \times p})$ and $z \perp\!\!\!\perp M(X)$. To apply the lower-tail bound for Gaussian quadratic forms, note that Hölder's inequality implies $\|M(X)\|_F \leq \sqrt{\text{tr}(M(X))}$ since $\|M(X)\|_{\text{op}} \leq 1$ almost surely. Therefore, letting $t = t' \text{tr}(M(X))$ with $t' \in (0, 1)$ in inequality (41) gives

$$\mathbb{P}\left[\Pi_p(z)^\top M(X) \Pi_p(z) \geq (1 - 2\sqrt{t'}) \text{tr}(M(X)) \mid M(X)\right] \geq 1 - \exp(-t' \text{tr}(M(X))) \quad (90)$$

Next, observe that inequality (41) with $t = t' \cdot n$ for $t' \in (0, 1)$ gives,

$$\mathbb{P}\left[\|z\|_2^2 \leq (1 + 4\sqrt{t'})n\right] \geq 1 - \exp(-t'n). \quad (91)$$

If we define the event

$$\mathcal{E}_1 := \left\{ \frac{1}{\frac{1}{\|z\|_2^2} \Pi_p(z)^\top M(X) \Pi_p(z)} \leq \frac{1+4\sqrt{t'}}{(1-2\sqrt{t'}) \text{tr}(M(X))} \frac{n}{\text{tr}(M(X))} \right\} \quad (92)$$

then the previous two inequalities imply

$$\begin{aligned} \mathbb{P}[\mathcal{E}_1 \mid M(X)] &\geq 1 - \exp(-t' \text{tr}(M(X))) - \exp(-t' \cdot n) \\ &\geq 1 - 2 \exp(-t' \text{tr}(M(X))), \end{aligned} \quad (93)$$

since $\text{tr}(M(X)) \leq n$ almost surely. Next, let $\kappa_5, \kappa_6 > 0$ be as in the previous lemma, and define the event

$$\mathcal{E}_2 := \left\{ \frac{n}{\text{tr}(M(X))} \leq \frac{1}{\kappa_5} n^{1-\frac{\gamma}{\eta}} \right\}, \quad (94)$$

which has probability $\mathbb{P}(\mathcal{E}_2) \geq 1 - \exp(-\kappa_6 n^{1+\frac{\gamma(1-\eta)}{\eta}})$.

We now put these items together. Starting with line (89), if we work on the intersection of \mathcal{E}_1 and \mathcal{E}_2 , then for any fixed $i = 1, \dots, n$ we have

$$\begin{aligned} \mathbb{P}\left[\frac{1}{v_\rho(X; X_i)} \leq \frac{1+4\sqrt{t'}}{(1-2\sqrt{t'})} \frac{1}{\kappa_5} n^{1-\frac{\gamma}{\eta}}\right] &\geq \mathbb{E}[1_{\mathcal{E}_1} \cdot 1_{\mathcal{E}_2}] \\ &\geq 1 - \mathbb{E}[1_{\mathcal{E}_1^c} + 1_{\mathcal{E}_2^c}] \\ &= 1 - \mathbb{E}[\mathbb{E}[1_{\mathcal{E}_1^c} \mid M(X)]] - \mathbb{P}(\mathcal{E}_2^c) \\ &\geq 1 - 2\mathbb{E}[\exp(-t' \text{tr}(M(X)))] - \mathbb{P}(\mathcal{E}_2^c) \\ &= 1 - 2\mathbb{E}[\exp(-t' \text{tr}(M(X))) \cdot (1_{\mathcal{E}_2} + 1_{\mathcal{E}_2^c})] - \mathbb{P}(\mathcal{E}_2^c) \\ &\geq 1 - \exp(-t' \kappa_5 n^{\frac{\gamma}{\eta}}) - 3\mathbb{P}(\mathcal{E}_2^c) \\ &\geq 1 - \exp(-t' \kappa_5 n^{\frac{\gamma}{\eta}}) - 3 \exp(-\kappa_6 n^{1+\frac{\gamma(1-\eta)}{\eta}}) \\ &\geq 1 - 4 \exp(-\min\{t' \kappa_5, \kappa_6\} \cdot n^{\frac{\gamma}{\eta}}) \end{aligned} \quad (95)$$

where we have used the previous lemma to bound $\mathbb{P}(\mathcal{E}_2^c)$, and also the assumption $\gamma \in (0, 1)$ to conclude that $\frac{\gamma}{\eta} \leq 1 + \frac{\gamma(1-\eta)}{\eta}$. Taking a union bound over $i = 1, \dots, n$, proves the claim. \square

The last component of Lemma 2 is to prove the MSPE inequalities.

D.3 Proof of the MSPE inequalities

The proof of Theorem (3) shows that for any realization of X we have

$$\text{mspe}(\hat{\beta}_\varrho | X) := \frac{1}{n} \mathbb{E} \left[\|X(\hat{\beta}_\varrho - \beta)\|_2^2 \mid X \right] \lesssim n^{-\theta} \|\beta\|_2^2 + \frac{1}{n} \sum_{i=1}^{p \wedge n} \left(\frac{l_i}{l_i + \varrho} \right)^2, \quad (96)$$

where $l_i = \lambda_i(\frac{1}{n} X^\top X)$. Now observe that the second term on the right side matches the expression for $\text{tr}(M(X))$ given in line (63) by replacing ρ with ϱ and multiplying by a factor of $\frac{1}{n}$. Therefore, using Lemma 10 and recalling $\varrho/n = n^{-\theta}$ shows that there are absolute constants $\kappa_1, \kappa_2, \kappa_3, \kappa_4 > 0$ such that for all large n ,

$$\mathbb{P} \left[\text{mspe}(\hat{\beta}_\varrho | X) \geq \kappa_1 (n^{-\theta} + n^{2(\theta-\eta)}) \right] \leq \exp(-\kappa_2 n^{2(1-\eta)+\theta}), \quad \text{if } \eta \in (0, \frac{1}{2}). \quad (97)$$

and

$$\mathbb{P} \left[\text{mspe}(\hat{\beta}_\varrho | X) \geq \kappa_3 (n^{-\theta} + n^{\frac{\theta}{\eta}-1}) \right] \leq \exp(-\kappa_4 n^{1+\frac{\theta(1-\eta)}{\eta}}), \quad \text{if } \eta > \frac{1}{2}. \quad (98)$$

In line (97), the bound $\text{mspe}(\hat{\beta}_\varrho | X)$ is optimized when $n^{-\theta} \asymp n^{2(\theta-\eta)}$, which explains the choice $\theta = \frac{2\eta}{3}$. Similarly, in line (98), the bound is optimized when $n^{-\theta} \asymp n^{\frac{\theta}{\eta}-1}$, which explains the choice $\theta = \frac{\eta}{\eta+1}$. Substituting in these values θ yields the stated result. \square

E Background results

E.1 Results on matrices and convexity

Lemma 12. Let $M_1 \in \mathbb{R}^{k_1 \times k_1}$ and $M_2 \in \mathbb{R}^{k_1 \times k_2}$. Then,

$$\|M_1 M_2\|_F \leq \|M_1\|_{\text{op}} \|M_2\|_F. \quad (99)$$

Proof. Observe that

$$\begin{aligned} \|M_1 M_2\|_F^2 &= \text{tr}(M_2^\top M_1^\top M_1 M_2) \\ &= \text{tr}((M_1^\top M_1)(M_2 M_2^\top)) \\ &\leq \sum_{i=1}^{k_1} \lambda_i(M_1^\top M_1) \cdot \lambda_i(M_2 M_2^\top) \\ &\leq \|M_1\|_{\text{op}}^2 \sum_{i=1}^{k_1} \lambda_i(M_2 M_2^\top) \\ &= \|M_1\|_{\text{op}}^2 \|M_2\|_F^2, \end{aligned} \quad (100)$$

where we have used von Neumann's trace inequality (also known as Fan's inequality) [23, p.10] in the third line. \square

A result on convex trace functionals. In the following lemma, an interval of the real line refers to any set of the form $(a, b), (a, b], [a, b)$, or $[a, b]$, where $-\infty \leq a \leq b \leq \infty$. We also define $\text{spec}(M)$ to be the set of eigenvalues of a square matrix M . The collection of symmetric matrices in $\mathbb{R}^{p \times p}$ is denoted by $\mathbb{S}^{p \times p}$. For a univariate function φ , the symbol $\text{tr}(\varphi(M))$ denotes $\sum_i \varphi(\lambda_i(M))$.

Lemma 13. Let $\mathcal{I} \subset \mathbb{R}$ be an interval, and let $\mathcal{M} \subset \mathbb{S}^{p \times p}$ be a convex set such that $\text{spec}(M) \subset \mathcal{I}$ for all $M \in \mathcal{M}$. Let $\varphi : \mathcal{I} \rightarrow \mathbb{R}$ be a convex function. Then, the functional

$$M \mapsto \text{tr}(\varphi(M)) \quad (101)$$

is convex on \mathcal{M} .

A proof may be found in the paper [24, Proposition 2].

Operator Jensen inequality. A function $f : \mathbb{S}^{p \times p} \rightarrow \mathbb{S}^{p \times p}$ is said to be operator convex if for all $\lambda \in [0, 1]$, and all $A, B \in \mathbb{S}^{p \times p}$,

$$f(\lambda A + (1 - \lambda)B) \preceq \lambda f(A) + (1 - \lambda)f(B), \quad (102)$$

where $A \preceq B$ means that $B - A$ is positive semidefinite.

Lemma 14 (Operator Jensen inequality). *Suppose $f : \mathbb{S}^{p \times p} \rightarrow \mathbb{S}^{p \times p}$ is operator convex, and let A be a random $\mathbb{S}^{p \times p}$ -valued matrix that is integrable. Then,*

$$f(\mathbb{E}[A]) \preceq \mathbb{E}[f(A)]. \quad (103)$$

Proof. It is enough to show that for all $x \in \mathbb{R}^p$,

$$x^\top f(\mathbb{E}[A])x \leq x^\top \mathbb{E}[f(A)]x. \quad (104)$$

For any fixed x , consider the function $g : \mathbb{S}^{p \times p} \rightarrow \mathbb{R}$ defined by $g(A) = x^\top f(A)x$. It is clear that g is a convex function in the usual sense, and so the ordinary version of Jensen's inequality implies $g(\mathbb{E}[A]) \leq \mathbb{E}[g(A)]$, which is the same as (104). \square

E.2 Results on Gaussian vectors and matrices

The following lemma is standard and is often referred to as the Gaussian concentration inequality [25].

Lemma 15. *Let $Z \in \mathbb{R}^p$ be a standard Gaussian vector and let $f : \mathbb{R}^p \rightarrow \mathbb{R}$ be an L -Lipschitz function with respect to the ℓ_2 norm. Then for all $t > 0$,*

$$\mathbb{P}\left(|f(Z) - \mathbb{E}[f(Z)]| \geq t\right) \leq 2 \exp\left(-\frac{t^2}{2L^2}\right). \quad (105)$$

Next, we give a formula for the expected square of a Wishart matrix.

Lemma 16. *Let $X \in \mathbb{R}^{n \times p}$ have rows drawn i.i.d. from $N(0, \Sigma)$, and let $\hat{\Sigma} = \frac{1}{n} X^\top X$. Then,*

$$\mathbb{E}[\hat{\Sigma}^2] = (1 + \frac{1}{n})\Sigma^2 + \frac{\text{tr}(\Sigma)}{n}\Sigma.$$

Proof. Write $\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n X_i X_i^\top$ where $X_i^\top \in \mathbb{R}^p$ is the i th row of X . If we put $M_i = X_i X_i^\top$, then

$$\hat{\Sigma}^2 = \frac{1}{n^2} \sum_{i=1}^n M_i^2 + \frac{1}{n^2} \sum_{i \neq j} M_i M_j.$$

Clearly, the M_i are independent with $\mathbb{E}[M_i] = \Sigma$ for all i , and so

$$\mathbb{E}[\hat{\Sigma}^2] = \frac{1}{n} \mathbb{E}[M_1^2] + \frac{2}{n^2} \binom{n}{2} \Sigma^2.$$

It remains to compute $\mathbb{E}[M_1^2]$. Write $X_i = \Sigma^{1/2} Z_i$ where Z_i is a standard Gaussian vector in \mathbb{R}^p , and let $\Sigma^{1/2} = U \Lambda^{1/2} U^\top$ be a spectral decomposition of $\Sigma^{1/2}$ where $U \in \mathbb{R}^{p \times p}$ is orthogonal, and Λ is diagonal with $\Lambda_{ii} = \lambda_i(\Sigma)$. By the orthogonal invariance of the normal distribution, $X_i \stackrel{\mathcal{L}}{=} U \Lambda^{1/2} Z_i$, and so

$$\mathbb{E}[M_1^2] = \mathbb{E}[X_1 X_1^\top X_1 X_1^\top] = U \Lambda^{1/2} \mathbb{E}[Z_1 Z_1^\top \Lambda Z_1 Z_1^\top] \Lambda^{1/2} U^\top. \quad (106)$$

Define the matrix $\tilde{M}_1 := Z_1 Z_1^\top \Lambda Z_1 Z_1^\top$. It is straightforward to verify that $\mathbb{E}[\tilde{M}_1]$ is diagonal, and its j th diagonal entry is

$$\mathbb{E}[\tilde{M}_1]_{jj} = \text{tr}(\Sigma) + 2\lambda_j.$$

Therefore,

$$\mathbb{E}[\tilde{M}_1] = \text{tr}(\Sigma) I_{p \times p} + 2\Lambda,$$

and combining this with line (106) gives

$$\mathbb{E}[M_1^2] = \text{tr}(\Sigma)\Sigma + 2\Sigma^2.$$

\square

Signed s.v.d. The following lemma describes the factors of an s.v.d. of a standard Gaussian matrix $Z \in \mathbb{R}^{n \times p}$ with $n \geq p$. To make the statement of the lemma more concise, we define the term *signed s.v.d.* below. This is merely a particular form of the s.v.d. that ensures uniqueness. Specifically, if $Z \in \mathbb{R}^{n \times p}$ is a full rank matrix with $n \geq p$, then the signed s.v.d. of Z is given by

$$Z = HLG^\top, \quad (107)$$

where $H \in \mathbb{R}^{n \times p}$ has orthonormal columns, the matrix $L \in \mathbb{R}^{p \times p}$ is diagonal with $L_{11} \geq L_{22} \geq \dots \geq L_{pp} > 0$, and $G \in \mathbb{R}^{p \times p}$ is orthogonal with its first row non-negative, i.e. $G_{1i} \geq 0$ for all $i = 1, \dots, p$. It is a basic fact from linear algebra that the signed s.v.d. of any full rank matrix in $\mathbb{R}^{n \times p}$ exists and is unique [21, Lemma 7.3.1]. This fact applies to Gaussian matrices, since they are full rank with probability 1.

Lemma 17. Suppose $n \geq p$, and let $Z \in \mathbb{R}^{n \times p}$ be a random matrix with entries drawn i.i.d. from $N(0, 1)$. Let

$$Z = HLG^\top, \quad (108)$$

be the unique signed s.v.d. of Z as defined above. Then, the matrices H , L and G are independent. Furthermore, if $H_i^\top \in \mathbb{R}^p$ denotes the i th row of H , then for each $i = 1, \dots, n$, the marginal law of H_i^\top is given by

$$H_i^\top \stackrel{\mathcal{L}}{=} \frac{1}{\|z\|_2} \cdot \Pi_p(z), \quad (109)$$

where $z \in \mathbb{R}^n$ is a standard Gaussian vector, and $\Pi_p(z)$ is the projection operator onto the first p coordinates, i.e. $\Pi_p(z) = (z_1, \dots, z_p)$.

Proof. We first argue that H , L , and G are independent, and then derive the representation for H_i^\top in the latter portion of the proof.

Due to the fact that the transformation $Z \mapsto (H, L, G)$ is invertible, it is possible to obtain the joint density of (H, L, G) from the density of Z by computing the matrix Jacobian of the factorization $Z = HLG$. (See the references [22], [26], and [27] for more background on Jacobians of matrix factorizations.) To speak in more detail about the joint density, let $\mathbb{V}^{n \times p}$ denote the Stiefel manifold of $n \times p$ matrices with orthonormal columns. Also, let $\mathbb{D}^{p \times p}$ denote the set of $p \times p$ diagonal matrices, and $\mathbb{O}^{p \times p}$ the set of orthogonal $p \times p$ matrices. The subset of $\mathbb{O}^{p \times p}$ with non-negative entries in the first row will be denoted by $\mathbb{O}_+^{p \times p}$.

Let $f_{H,L,G} : \mathbb{V}^{n \times p} \times \mathbb{D}^{p \times p} \times \mathbb{O}_+^{p \times p} \rightarrow [0, \infty)$ denote the joint density of (H, L, G) , where the base measure is the product of Haar measure on $\mathbb{V}^{n \times p}$, Lebesgue measure on $\mathbb{D}^{p \times p}$, and Haar measure on $\mathbb{O}^{p \times p}$ restricted to $\mathbb{O}_+^{p \times p}$. (See the book [28] for background on these measures). From lines 8.8-8.10 in the paper [29], it is known that such a density exists, and that $f_{H,L,G}$ factors according to

$$f_{H,L,G}(h, l, g) = f_H(h) \cdot f_L(l) \cdot f_G(g), \quad (110)$$

where the function f_H denotes the density of H with respect to Haar measure on $\mathbb{V}^{n \times p}$, the function f_L denotes the density of L with respect to Lebesgue measure on $\mathbb{D}^{p \times p}$, and the function f_G denotes the density of G with respect to the restriction of Haar measure to $\mathbb{O}_+^{p \times p}$. This proves that H , L , and G are independent.

We now prove the representation (109). From line 8.10 in the paper [29], it is known that f_H is constant on $\mathbb{V}^{n \times p}$. In other words, the matrix H follows the Haar distribution on $\mathbb{V}^{n \times p}$. Consequently, Theorem 2.2.1(ii) in the book [28] implies that the rows H_i^\top can be represented as

$$H_i^\top \stackrel{\mathcal{L}}{=} \Pi_p(J_i^\top) \quad (111)$$

where $J_i^\top \in \mathbb{R}^n$ is the i th row of a Haar-distributed random matrix in $\mathbb{O}^{n \times n}$. Furthermore, the rows J_i^\top are uniformly distributed on the unit sphere on \mathbb{R}^n , and hence can be represented as $z/\|z\|_2$, where $z \in \mathbb{R}^n$ is a standard Gaussian vector. \square

References

- [1] C.-H. Zhang and S. S. Zhang. Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society: Series B*, 76(1):217–242, 2014.
- [2] A. Javanmard and A. Montanari. Hypothesis testing in high-dimensional regression under the Gaussian random design model: Asymptotic theory. *arXiv preprint arXiv:1301.4240*, 2013.
- [3] A. Javanmard and A. Montanari. Confidence intervals and hypothesis testing for high-dimensional regression. *arXiv preprint arXiv:1306.3171*, 2013.
- [4] P. Bühlmann. Statistical significance in high-dimensional linear models. *Bernoulli*, 19(4):1212–1242, 2013.
- [5] S. van de Geer, P. Bühlmann, and Y. Ritov. On asymptotically optimal confidence regions and tests for high-dimensional models. *arXiv preprint arXiv:1303.0518*, 2013.
- [6] J. D. Lee, D. L. Sun, Y. Sun, and J. E. Taylor. Exact inference after model selection via the lasso. *arXiv preprint arXiv:1311.6238*, 2013.
- [7] A. Chatterjee and S. N. Lahiri. Rates of convergence of the adaptive lasso estimators to the oracle distribution and higher order refinements by the bootstrap. *The Annals of Statistics*, 41(3):1232–1259, 2013.
- [8] H. Liu and B. Yu. Asymptotic properties of lasso+mls and lasso+ridge in sparse high-dimensional linear regression. *Electronic Journal of Statistics*, 7:3124–3169, 2013.
- [9] V. Chernozhukov, D. Chetverikov, and K. Kato. Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *The Annals of Statistics*, 41(6):2786–2819, 2013.
- [10] E. L. Lehmann and J. P. Romano. *Testing Statistical Hypotheses*. Springer, 2005.
- [11] D. A. Freedman. Bootstrapping regression models. *The Annals of Statistics*, 9(6):1218–1228, 1981.
- [12] P. J. Bickel and D. A. Freedman. Bootstrapping regression models with many parameters. In *Festschrift for Erich L. Lehmann*, pages 28–48. Wadsworth, 1983.
- [13] N. R. Draper and H. Smith. *Applied Regression Analysis*. Wiley-Interscience, 1998.
- [14] P. J. Bickel and D. A. Freedman. Some asymptotic theory for the bootstrap. *The Annals of Statistics*, pages 1196–1217, 1981.
- [15] S. Bobkov and M. Ledoux. *One-dimensional empirical measures, order statistics, and Kantorovich transport distances*. preprint, 2014.
- [16] A. B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2009.
- [17] L. Wasserman. *All of Nonparametric Statistics*. Springer, 2006.
- [18] B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, 28(5):1302–1338, 2000.
- [19] R. A. Horn and C. R. Johnson. *Topics in matrix analysis*. Cambridge University Press, 1991.
- [20] R. Bhatia. *Matrix Analysis*. Springer, 1997.
- [21] R. A. Horn and C. A. Johnson. *Matrix Analysis*. Cambridge University Press, 22nd printing edition, 2009.
- [22] R. J. Muirhead. *Aspects of Multivariate Statistical Theory*. John Wiley & Sons, 1982.
- [23] J. M. Borwein and A. S. Lewis. *Convex Analysis and Nonlinear Optimization Theory and Examples*. CMS Books in Mathematics. Canadian Mathematical Society, 2000.

- [24] D. Petz. A survey of certain trace inequalities. *Functional analysis and operator theory*, 30:287–298, 1994.
- [25] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- [26] A. M. Mathai. *Jacobians of Matrix Transformations and Functions of Matrix Argument*. World Scientific, 1997.
- [27] A. Edelman and N. R. Rao. Random matrix theory. *Acta Numerica*, 14:233–297, 2005.
- [28] Y. Chikuse. *Statistics on Special Manifolds*, volume 174. Springer Science & Business Media, 2003.
- [29] A. T. James. Normal multivariate analysis and the orthogonal group. *The Annals of Mathematical Statistics*, pages 40–75, 1954.